arXiv:1603.02384v1 [math.OC] 8 Mar 2016

Column Generation and Lazy constraints for solving the Liner Ship Fleet Repositioning Problem with cargo flows

Robin H. Pearce, Alexis Tyler & Michael Forbes

School of Mathematics and Physics, University of Queensland, Australia

Abstract

We consider an important problem in the shipping industry known as the liner shipping fleet repositioning problem (LSFRP). We examine a public data set for this problem including many instances which have not previously been solved to optimality. We present several improvements on a previous mathematical formulation, however the largest instances still result in models too difficult to solve in reasonable time. The implementation of column generation reduces the model size significantly, allowing all instances to be solved, with some taking two to three hours. A novel application of lazy constraints further reduces the size of the model, and results in all instances being solved to optimality in under four minutes.

Keywords: Liner Shipping, Network flows, Scheduling, Column Generation, Lazy constraints

1. Introduction

The shipping industry is responsible for transporting many billion tons of cargo each year, almost 9.6 billion tons in 2013, with a large percentage of this (35.4%) made up by container shipping [24]. Optimising the practices of the shipping industry to maximise profits is the subject of considerable academic and commercial interest. Much of this research centres around the principle of optimising scheduling to maximise profits. The main areas of research are in the fleet deployment problem and the network design problem, which are related to this principle.

Liner shipping operates on a fixed schedule and uses a standard container size, the ISO container [20]. Within the shipping industry, ships are often referred to as vessels, and the set of ships which a liner shipping company operates is called a fleet. An operator is the person responsible for the operation of the shipping company and its fleet. These companies operate within a shipping network, which is a set of ports connected by arcs over which the fleet may travel. While liner

shipping does operate on a fixed, periodic schedule, ships within the fleet often need to be moved between services to meet the demand of clients. This often happens to create new or modify existing services. However, repositioning vessels is an expensive exercise due to fuel costs and potential losses in revenue, so optimising these journeys is of particular interest within the industry.

In this paper we present several improvements for solving the LSFRP. We begin with an improved MIP model for the LSFRP which ignores some less important aspects of the problem. Next, we apply column generation and present a proof of the integer properties of the column generation master problem. Finally, we present a novel use of lazy constraints to massively speed up the solution of the column generation formulation.

This paper is organised as follows: the remainder of this section will be a review of the literature around shipping network optimisation, followed by a review more specific to the LSFRP. In Section 2 we will describe the problem in more detail, and in Section 3 we will look at the improvements we have made to the original formulation. Section 4 contains computational results comparing our different formulations, and Section 5 is a discussion about various aspects of this problem and the associated public data set.

While the optimisation of network flow models is well studied, very little work focusses on the application of these models to shipping networks. This is emphasised in a survey conducted by Ronen [18] in which he states that, while the development of these models has the potential to have a large economic impact, the lack of literature in the field has meant that the application of these models in industry has been limited. Ronen details the models and techniques that were used in ship scheduling in the previous decade. He notes that the majority of the work was focused on four main areas: fleet size, mix and deployment; inventory routing; cruising speed; and ship scheduling. It is interesting to note that only a handful of the papers reviewed in this survey pertain to liner shipping.

One such study is by Brown, Graves and Ronen [5], and focusses on scheduling ocean transportation of crude oil. The authors use an elastic set partitioning model which incorporates all fleet cost components and optimises for speed as well as scheduling. However, this model assumes that all ships are of similar sizes, and that one unit of cargo is a full shipload which has only a single discharging port, which, in reality, are not feasible assumptions. With these assumptions the authors are able to solve this problem with thousands of binary variables to an optimal integer solution in less than a minute.

Another study is by Rana and Vickson [16], in which they use Lagrangian Relaxation to solve the container ship routing problem. Their model incorporates multiple ships and solves for the optimal schedule and amount of cargo transported by the fleet to maximise profit. The constraints in this model cover the carrying capacity of the ships, available cargo, time duration, and the requirement of having a connected feasible route for each ship. This model does not incorporate more complex scheduling constraints for the ships, for example a cargo delivery deadline, which limits the potential applications for this model.

In their 1992 study, Rathi, Church and Solanki [17] explore allocating resources to support a multi-commodity flow with time windows. This study presents three different linear programming (LP) models to minimise costs associated with this problem. All three formulations achieve this by minimising lateness, which is determined by the amount of assets of each type that should be allocated to a given route in each time period.

The review by Christiansen, Fagerholt and Ronen [7] provides an overview of the published research on ship routing and scheduling from 1994-2002. The review is split into several sections. The first section looks at strategic fleet planning, the design of fleets. The second section is tactical and operational fleet planning which explores ship routing and scheduling problems for different types of shipping. Fagerholt and Christiansen [9] solve a combined ship scheduling and allocation problem for industrial shipping using set partitioning. Next, they consider a problem related to robust ship scheduling with multiple time windows using a set partitioning method, where all feasible ship schedules are found prior to the model being solved [6]. Sherali [19] investigates fleet management models and algorithms for an oil-tanker routing and scheduling problem using a mixed integer program (MIP) with heuristics.

The third and final section of their review covers naval applications and other related problems. Of the papers reviewed pertaining to ship routing and scheduling, most used dynamic programming (DP), integer programs (IP) or heuristic algorithms to solve the problem. The authors note that trends in the literature display a growing need for research into these types of problems.

There are a number of studies dedicated to solving the fleet deployment problem, such as the one by Powell and Perkins [15], and the network design problem, for example those by Agarwal and Ergun [1], Álvarez [2] and Brouer et. al. [4]. These problems are related to the decisions associated with the design of networks and timetables, as opposed to determining optimal paths for ships through a predetermined network. Agarwal and Ergun provide a unique formulation of

ship scheduling and network design for cargo routing in liner shipping. The authors use a mixed integer linear program to solve the ship scheduling and cargo routing problem. Their model includes constraints such as frequency of operation and transshipment of cargo.

Agarwal and Ergun also develop several solution techniques and compare their success and efficiency. The described algorithms are a greedy heuristic, column generation and Benders decomposition. They also outline an iterative search algorithm to generate schedules for liner shipping. Computational tests were performed on each of the algorithms using generated instances and it was found that, while the greedy heuristic was able to solve the problem very quickly, the solution quality was low. It was also found that the column generation and Benders' decomposition approaches achieved similar results in terms of accuracy, however the Benders' formulation was significantly faster, particularly when the instance size was large.

1.1. Liner ship fleet repositioning problem

The LSFRP is a type of network design problem which involves repositioning ships between service routes while maximising profit. This is achieved by visiting ports and delivering cargo while repositioning. However, despite the body of literature devoted to liner shipping and its surrounding problems, very little of this research is focused on the liner shipping fleet repositioning problem.

The first study that explores the LSFRP is by Tierney et. al. [21]. In this paper, the authors solve a simplified version of the LSFRP without cargo flows, empty equipment, or sail-on-service (SOS) opportunities (discussed further below). Tierney and Jensen then conduct a study which continues to explore this problem and incorporated cargo flows [22]. In this paper, the authors use a mixed-integer program (MIP) in conjunction with a constructed graph to solve the LSFRP. This graph incorporates many of the LSFRP-specific constraints (such as SOS opportunities) so that they can be removed from the model formulation. This approach is able to solve several instances to optimality, however there are many larger instances where the problem can not be solved, as the solver runs out of memory or exceeds the maximum CPU time of one hour.

Another approach to solving the LSFRP is proposed by Kelareva, Tierney and Kilby [13]. In this study, they solve the full LSFRP with SOS opportunities, however they do not incorporate cargo flows into their model. A constraint programming (CP) method is used with lazy clause generation, and is tested against the MIP in [21]. After testing the different models on a data set, the CP method is found to be faster than the MIP for all instances. However, this only occurs after

choosing a search strategy for the particular problem. The authors note that, without sufficient understanding of the problem and of CP modelling techniques, it is difficult to choose a search strategy that is both fast and successfully finds optimal solutions. Furthermore, the CP method can not be extended to allow pre-computations or chaining of SOS or opportunities to carry empty cargo containers.

The most recent study on the LSFRP is by Tierney et. al. [20], which expands on the work by Tierney and Jensen from 2012 [22]. They improve the model, provide a public data set, and use a heuristic approach. This model was able to incorporate many complex aspects of the LSFRP, including SOS opportunities, phase-in/phase-out requirements, and flexible arcs. Some of these (SOS opportunities and phase-in/phase-out requirements) are processed into the graph structure, along with sailing costs and cabotage restrictions. The MIP forms a "disjoint path problem in which a fractional multi-commodity flow is allowed to flow over arcs in the vessel paths, along with a small scheduling component in the flexible nodes" [20].

2. Problem description and model formulation

The LSFRP consists of finding sequences of activities that move vessels between services in a liner shipping network, while maximising profit by trading off ship moving costs and cargo flow incomes [20]. "Liner shipping services are composed of multiple slots, each of which represents a cycle that is assigned to a particular vessel". The slots contain nodes or ports which must be visited by vessels at specific times in sequence. When a vessel is assigned a slot, it sails to all of its ports in order and delivers its cargo. Figure 1 shows an example of a service with three slots (represented by the differently styled lines in the graph) and five ports (a, b, c, d, e). The diagram shows that each slot takes three weeks to return to the start of the cycle, so three ships would be needed to run this service weekly. An LSFRP needs to be solved when we are transitioning a fleet from one set of services to a new one.

Another aspect of repositioning that needs to be taken into account is the time constraints. The time at which a ship may begin repositioning is known as the phase-out time. The ship must finish repositioning by the phase-in time of the goal service. In between these two times the ship is available for repositioning and is able to undertake a number of activities to both reach its goal service and reduce costs. An example of this can be seen in Figure 1. The latest phase in time is

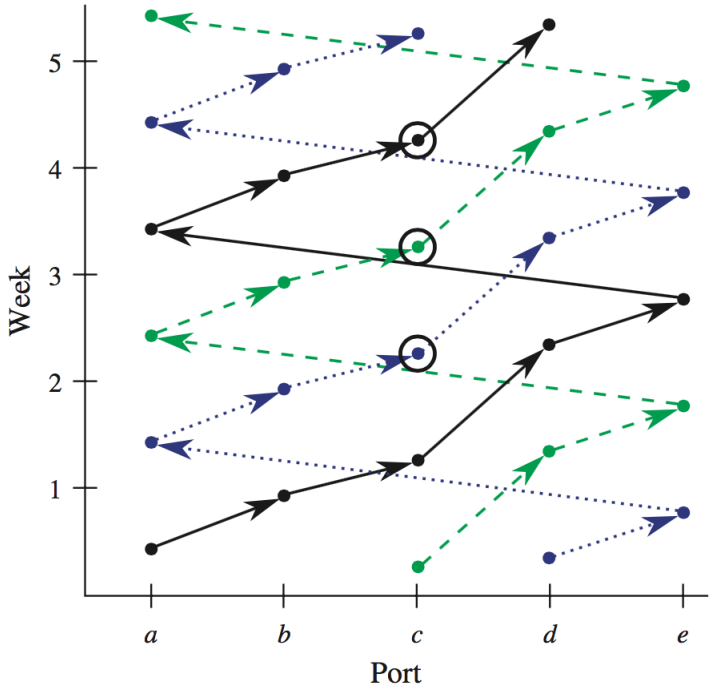

Figure 1: A time-space graph of a service with three vessels. Image sourced from Tierney et. al. (2014) [20]

at port c in week 2 and the phase out time is after all the hollow circles [20]. Between these points a vessel may undertake repositioning activities.

The LSFRP is best described using Figure 2. This shows a ship which needs to be repositioned from its initial service (Chennai Express) to a goal service (Intra-WCSA). During repositioning the vessel can deliver cargo to ports to offset the cost of moving the ship, thus cargo flows are an important aspect of the problem. One way to do this is to take advantage of sail-on-service (SOS) opportunities, which are situations in which a repositioning ship can replace an on-service vessel for part of its service in order to reduce costs (by not having two ships sailing on the same course unnecessarily). There are two main methods of performing a SOS opportunity: transhipping, where all cargo from the on-service ship is moved onto the repositioning ship at a port, or parallel sailing, where the two ships visit the same ports sequentially and the on-service vessel only unloads cargo, while the repositioning vessel only loads cargo.

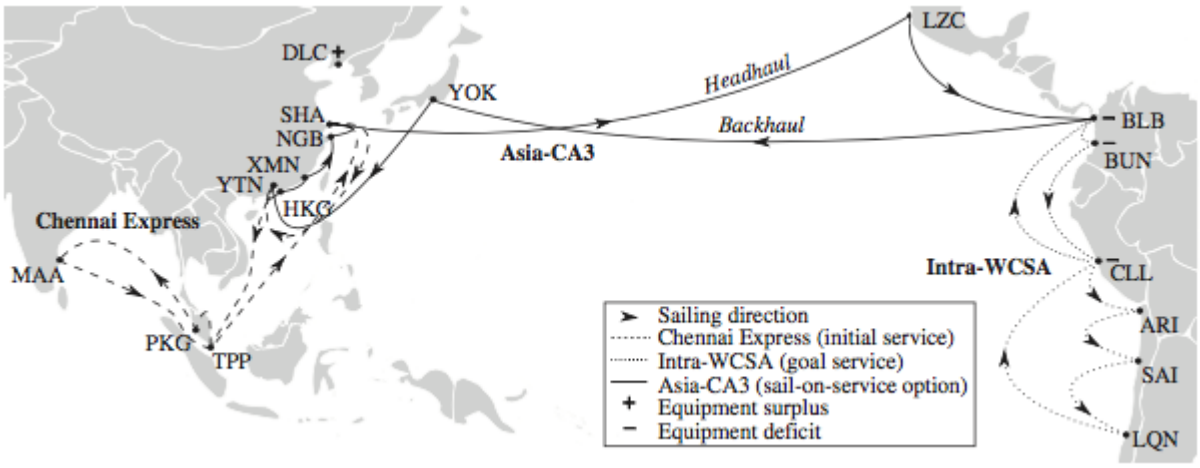

Figure 2: Liner shipping network. Image sourced from Tierney et. al. (2014) [20]

2.0.1. Empty containers and flexible arcs

Another way for ships to offset the cost of repositioning is to transport empty containers from ports with an empty equipment surplus to ports with a deficit. The revenue from performing this type of activity is calculated as an approximation of the savings from moving the equipment now, as opposed to at a later date, potentially through a more expensive channel. There are two types of cargo considered in this problem: dry and refrigerated (reefer). We must differentiate between the two types, since when transporting cargo the reefer containers must be plugged into a power outlet, which means that ships will only have a limited reefer capacity. This is not the case when moving empty equipment, however we still make the distinction as the deficit we are supplying may be for a specific container type. We use the term flexible visits to denote ports with empty equipment available, but no actual cargo demands. These flexible visits are travelled to via flexible arcs.

There are also various restrictions placed on the cargo carried by repositioning ships such as trade zones. Trade zones are countries or groups of countries with trade agreements. Often cargo cannot flow between trade zones without violating these agreements. To avoid the movement of cargo violating these trade zone restrictions, the law, or a customer contract, repositioning ships are disallowed from crossing into other trade zones while carrying cargo. A similar restriction is known as a cabotage restriction, which prevents international ships from performing domestic cargo services [20]. These are all aspects which need to be considered when modelling the LSFRP. Most

of these restrictions have been incorporated directly into the network of potential ship paths, so they will not be represented in the MIP formulation.

For the original model formulation, we refer the reader to the paper by Tierney et. al. (2014) [20]. We have maintained consistency in notation from previous studies. We now present a reduced formulation of the LSFRP.

2.1. Reduced MIP

Starting with the MIP model described in Tierney et. al. (2014) [20], a reduced version is formulated that does not incorporate flexible arcs or empty equipment. These aspects of the problem are omitted for simplicity to allow us to explore the core structure of the LSFRP without added complexity. By noting the percentage of the public data set that does not include these additional requirements (66%), it can be seen that the reduced problem is still able to provide much value, as the majority of the instances do not contain the more complex aspects of the problem. There are more reasons why omitting these aspects of the problem are reasonable, some of which will be explored in Section 5.

Parameters

S	Set of ships.
V'	Set of visits minus the graph sink.
A'	Set of arcs minus those arcs connecting to the graph sink, i.e., $(i,j) \in A, i,j \in V'$.
Q	Set of cargo types; $Q = \{dc, rf\}$.
M	Set of demand triplets of the form (o,d,q), where $o \in V'$, $d \subseteq V'$, and $q \in Q$ are the origin visit, possible destination visits, and the cargo type respectively.
$M_i^{\text{Orig}}, (M_i^{\text{Dest}}) \subseteq M$	Set of demands with an origin (destination) visit $i \in V$.
$u_s^q \in \mathbb{R}^+$	Capacity of vessel s for cargo type $q \in Q$.
$v_s \in V'$	Starting visit of ship $s \in S$.

$r^{(o,d,q)} \in \mathbb{R}^+$	Amount of revenue gained per TEU of loaded containers carried for the demand triplet.
$c_{sij}^{\text{Sail}} \in \mathbb{R}^+$	Fixed cost of vessel s utilizing are $(i,j) \in A'$.
$c_i^{\text{Mv}} \in \mathbb{R}^+$	Cost to move a single TEU on or off a ship at visit $i \in V'$.
$c_{si}^{\text{Port}} \in \mathbb{R}$	Port fee associated with vessel s at visit $i \in V'$.
$a^{(o,d,q)} \in \mathbb{R}^+$	Amount of demand available for the demand triplet.
$In(i) \subseteq V'$	Set of visits with an arc connecting to visit $i \in V$.
$Out(i) \subseteq V'$	Set of visits receiving an arc from visit $i \in V$.
$\tau \in V$	Graph sink, which is not an actual visit.

Variables

$x_{ij}^{(o,d,q)} \in \mathbb{R}_0^+$	Amount of flow of demand triplet $(o,d,q) \in M$ on $(i,j) \in A'$.
$y_{ij}^s \in \{0,1\}$	Indication of whether vessel s is sailing on arc $(i,j) \in A$.

Objective and Constraints

$$\max \quad \left\{ \sum_{(o,d,q)\in M} \left(\sum_{j\in d} \sum_{i\in In(j)} (r^{(o,d,q)} - c_o^{Mv} - c_j^{Mv}) x_{ij}^{(o,d,q)} \right) \right. \tag{1}$$

$$\left. - \sum_{s\in S} \sum_{(i,j)\in A'} c_{sij}^{\text{Sail}} y_{ij}^s - \sum_{j\in V'} \sum_{i\in In(j)} \sum_{s\in S} c_{sj}^{\text{Port}} y_{ij}^s \right\} \tag{2}$$

$$\text{s.t.} \quad \sum_{s\in S} \sum_{i\in In(j)} y_{ij}^s \leq 1, \quad \forall j \in V'; \tag{3}$$

$$\sum_{j\in Out(i)} y_{ij}^s = 1, \quad \forall s \in S, i = v_s; \tag{4}$$

$$\sum_{i\in In(\tau)} \sum_{s\in S} y_{i\tau}^s = |S|; \tag{5}$$

$$\sum_{i\in In(j)} y_{ij}^s - \sum_{i\in Out(j)} y_{ij}^s = 0, \quad \forall j \in V' \setminus \bigcup_{s\in S} v_s, s \in S; \tag{6}$$

$$\sum_{(o,d,rf)\in M} x_{ij}^{(o,d,rf)} \leq \sum_{s\in S} u_s^{rf} y_{ij}^s, \quad \forall (i,j) \in A'; \tag{7}$$

$$\sum_{(o,d,q)\in M} x_{ij}^{(o,d,q)} \leq \sum_{s\in S} u_s^{dc} y_{ij}^s, \quad \forall (i,j) \in A'; \tag{8}$$

$$\sum_{i\in Out(o)} x_{oi}^{(o,d,q)} \leq a^{(o,d,q)} \sum_{i\in Out(o)} \sum_{s\in S} y_{oi}^s, \quad \forall (o,d,q) \in M; \tag{9}$$

$$\sum_{i\in In(j)} x_{ij}^{(o,d,q)} - \sum_{k\in Out(j)} x_{jk}^{(o,d,q)} = 0, \quad \forall (o,d,q) \in M, j \in V'\backslash(o \cup d); \tag{10}$$

The objective function maximises the profit of the shipping company. The first line (1) calculates the profit from delivering the cargo by adding the revenue minus the cost to transport the cargo on and off the ship. This is multiplied by the amount of cargo carried. The second line of the objective function (2) subtracts the sum of the sailing costs and the port fees for each port visited by each ship.

Constraint (3) ensures that only one ship visits each port, while (4-6) conserve the flow of each ship from its starting port to the sink node. If a ship uses an arc, that arc is assigned a reefer capacity in (7) and a total capacity in (8). Constraint (9) ensures that cargo can only flow along an arc if it is on a ship. Constraint (10) conserves the flow of cargo from its source node to its destination by ensuring that if it enters an intermediate node, it must also exit that node.

The reduced MIP only uses the $x_{ij}^{(o,d,q)}$ and y_{ij}^s variables, as the others pertain to flexible arcs, empty equipment or entrance/exit times. As such, constraints and terms referring to the other variables are omitted from the reduced formulation. With the use of the Gurobi solver package [10] this model was implemented and solved using a subset of the public data instances (those which did not contain flexible arcs). Pre-processing ensures that only the $x_{ij}^{(o,d,q)}$ variables which can be non-zero are added to the model. That is, there exists a path from the origin to one of the demand points passing through arc (i,j).

3. Improved solution techniques for LSFRP

We have made a number of improvements to the solution of the LSFRP, which we describe in this section. The reduced MIP is our starting point, which eliminates some complexity while preserving the core components of the problem. We then tighten some constraints and reformulate the model using disaggregation, which increases the number of variables we solve for, but results in a tighter bound. Next, we apply column generation, which allows us to solve all previously unsolved instances within a large time window. Finally, we apply lazy constraints, which greatly simplifies the problem, to the point where all instances are solved to optimality within four minutes.

3.1. Tighter Bound and revised formulation

While the reduced MIP is able to replicate the results shown in Tierney et. al. (2014) [20], it is still unable to solve the last seven instances in the public data set. We note that one of the reasons the MIP struggles on larger problems is because the linear relaxation of the problem generates solutions in which fractional ship variables are used to transport all of a demand triplet. In order to prevent this an additional set of constraints is added:

$$x_{ij}^{(o,d,q)} \leq a^{(o,d,q)} \sum_{s \in S} y_{ij}^{s} \quad \forall (i,j) \in A', (o,d,q) \in M; \tag{9a}$$

These constraints prevent ships from moving a greater fraction of the demand triplet than the fraction of the ship used. This is a disaggregated version of constraint (9) from the reduced formulation, as it is no longer summed over $i \in Out(o)$. This is allowed since only one ship can visit any node, and thus only one arc leaving each node will have a non-zero value of y_{ij}^{s} in any integer solution. By the properties of disaggregation this must give a tighter bound for the linear relaxation. This improved bound yielded strong improvement on some larger instances, however it is still unable to solve five instances to optimality within the timeout limit.

Despite these tighter constraints, fractional parts of demand triplets can still be shipped, as the new constraints apply to all the ships rather than individual ships (the RHS is summed over s). To combat this, we reformulate the model for individual ships by adding a ship index to the x variables. As stated earlier, an important aspect of the problem to note is that the paths need to be node distinct, meaning that only one ship can visit each node. This property means that no product can be transshipped, and allows the ship index to be added to the x variables. The revised formulation with $x_{ij}^{s,(o,d,q)}$ variables is shown below.

Variables

$x_{ij}^{s,(o,d,q)} \in \mathbb{R}_0^+$ Amount of flow of demand triplet $(o,d,q) \in M$ on $(i,j) \in A'$ on $s \in S$.

$y_{ij}^{s} \in \{0,1\}$ Indication of whether vessel s is sailing on arc $(i,j) \in A$.

Objective and Constraints

$$\max \quad \left\{ \sum_{(o,d,q) \in M} \sum_{s \in S} \sum_{j \in d} \sum_{i \in In(j)} (r^{(o,d,q)} - c_o^{Mv} - c_j^{Mv}) x_{ij}^{s,(o,d,q)} \right. \tag{11}$$

$$-\sum_{s\in S}\sum_{(i,j)\in A'} c^{\text{Sail}}_{sij} y^s_{ij} - \sum_{j\in V'}\sum_{i\in In(j)}\sum_{s\in S} c^{\text{Port}}_{sj} y^s_{ij}\Bigg\} \tag{12}$$

s.t. $$\sum_{s\in S}\sum_{i\in In(j)} y^s_{ij} \le 1, \quad \forall j \in V'; \tag{13}$$

$$\sum_{j\in Out(i)} y^s_{ij} = 1, \quad \forall s \in S, i = v_s; \tag{14}$$

$$\sum_{i\in In(\tau)}\sum_{s\in S} y^s_{i\tau} = |S|; \tag{15}$$

$$\sum_{i\in In(j)} y^s_{ij} - \sum_{i\in Out(j)} y^s_{ji} = 0, \quad \forall j \in V'\backslash \bigcup_{s\in S} v_s, s \in S; \tag{16}$$

$$\sum_{(o,d,rf)\in M} x^{s,(o,d,rf)}_{ij} \le u^{rf}_s y^s_{ij}, \quad \forall (i,j) \in A', s \in S; \tag{17}$$

$$\sum_{(o,d,q)\in M} x^{s,(o,d,q)}_{ij} \le u^{dc}_s y^s_{ij}, \quad \forall (i,j) \in A', s \in S; \tag{18}$$

$$\sum_{i\in Out(o)} x^{s,(o,d,q)}_{oi} \le a^{(o,d,q)} \sum_{i\in Out(o)} y^s_{oi}, \quad \forall (o,d,q) \in M, s \in S; \tag{19}$$

$$\sum_{i\in In(j)} x^{s,(o,d,q)}_{ij} - \sum_{k\in Out(j)} x^{s,(o,d,q)}_{jk} = 0, \quad \forall (o,d,q) \in M, j \in V'\backslash(o \cup d), s \in S; \tag{20}$$

$$x^{s,(o,d,q)}_{ij} \le y^s_{ij} \min(a^{(o,d,q)}, u^q_s), \quad \forall (i,j) \in A', s \in S, (o,d,q) \in M; \tag{21}$$

The objective value is unchanged from the reduced MIP, except that now the x variables are also summed over all ships $s \in S$. Constraints (13-16) are identical to the original formulation, and constraints (17-20) are disaggregated versions of constraints (7-10), so there is now one constraint for each ship. Finally, constraint (21) is a disaggregated version of constraint (9a), which ensures that for each ship, and on each arc, no more cargo can be transported than is either available or able to be transported on the ship.

While this formulation does introduce more variables into the problem, it also provides a linear relaxation with a tighter bound, which allows it to solve much faster for larger instances. The results of this new MIP formulation are reported in Section 4.2.

3.2. Column Generation

One of the main reasons that larger instances are unsolvable when using the presented MIP formulations is because the problems are too large to consider all variables explicitly and still be solved within a reasonable time frame. The last three instances are so large that the LP relaxation can not be solved using any of the MIP formulations within one hour. To combat this issue,

composite variables and the technique of column generation are employed to reduce the number of variables considered at one time [3].

First the MIP is decomposed into a master problem and sub-problems (one for each ship). The master problem is an IP with every possible composite variable, which represent paths of ships through the network. This is then transformed into the reduced master problem (RMP), which contains a reduced set of the composite variables. We begin by solving the LP relaxation of the RMP, which is an LP that contains two constraints. The first says each ship is used exactly once, and the second says each node is used at most once. The variables for the master problem are paths through the graph for certain ships. The sub-problems are equivalent to the revised MIP on a ship-by-ship basis, with the dual variables associated with the master problem in the objective function. Since the path of each ship is decided by the master problem, the constraints on each node being visited only once are removed from the sub-problems. Since there is one sub-problem per ship, we do not need to sum over the ships. The RMP is described as:

Parameters.

P	Set of paths.
$C_{sp} \in \mathbb{R}^+$	Profit of vessel s sailing on path p (revenue from moving product less the cost of the path).
$\delta_{isp} \in \{0,1\}$	1 if vessel s sailing on path p goes through node $i \in V'$.

Variables.

$Z_{sp} \in \{0,1\}$ 1 if vessel s sails on path $p \in P$, 0 otherwise.

Objective and Constraints.

$$\max \quad \sum_{p \in P} \sum_{s \in S} C_{sp} Z_{sp} \tag{22}$$

$$\text{s.t.} \quad \sum_{p \in P} Z_{sp} = 1, \quad \forall s \in S \tag{23}$$

$$\sum_{p \in P} \sum_{s \in S} \delta_{isp} Z_{sp} \leq 1, \quad \forall i \in V' \tag{24}$$

After setting up the models, an initial solution to the master problem is generated. This is achieved using a modified version of the sub-problems. First the number of possible paths through

the graph for each ship is calculated, then the ships are ordered from those with the least to those with the most paths. The sub-problems are solved in this order using the objective function from the revised formulation, with s as a constant. After each subproblem is solved, the nodes used in the solution are excluded from subsequent sub-problems. This heuristic procedure resulted in a feasible starting solution for each instance, though our code does allow for high cost dummy paths direct from source to sink if this procedure fails for any ship. The paths generated from these sub-problems are then added to the master problem as columns. The master problem is then solved using these columns.

In the main loop of the algorithm, the sub-problems are solved in the reverse order than before, and if a new column is found, it is added to the master problem, which is then re-solved. If all sub-problems are solved without any new column being added to the master problem, the algorithm terminates and the optimal solution has been found. Due to specific properties of this problem, the column generation master problem gives integer optimal solutions, so a branch and bound algorithm is not required. The reasons for this are explored in Section 5.1.

3.3. Lazy Constraints

The final improvement we present for the solution of the LSFRP involves using lazy constraints to dramatically reduce the size of the column generation sub-problems. Each node can be visited by at most one ship, the ship movement graph is acyclic, and no transshipment can occur. Therefore, if any two demands are ever carried together at any time, then they will always be carried together. This is because each demand has a single port of origin, and that port cannot be visited twice.

Consider a ship that visits port A and loads some demand from that port. The ship then proceeds to port B where it loads another demand from this port. Finally, the ship then proceeds to port C. Because the demands from ports A and B were carried together, the sum of demands from A and B cannot exceed the capacity of the ship on this arc. If demand B is unloaded at port C, it will not be possible to return to port A to load more demand of type A, and vice versa.

This means that rather than solving for the flow along each arc, we can instead solve for the total flow of each demand, and add extra constraints enforcing shared capacity where it is exceeded. We define a new set $\bar{M}_s$, which is the set of all demands (o, d, q) which can be moved by ship s. By extension, $\bar{V}_{sq}^{\text{Orig}}$ is the set of all nodes from which ship s can pick up a demand of type q from $\bar{M}_s$, that is:

$$\bar{V}_{sq}^{\text{Orig}} = \{o|(o,d,q) \in \bar{M}_s\}$$

Finally, we denote the set of all arcs $(i,j) \in A'$ across which a demand triple $(o,d,q) \in M$ can possibly travel as $A^{(o,d,q)}$. We also use the set M_i^{Orig} from the original formulation, which is the set of demands with an origin visit $i \in V$. Starting with the revised formulation from Section 3.1, we modify the x variables so that they are no longer indexed by arc. That is:

$$x_s^{(o,d,q)} \in \mathbb{R}_0^+ \quad \text{Amount of flow of demand triplet } (o,d,q) \in M \text{ on ship } s \in S.$$

We remove constraints (17-21) and replace them with the following:

$$\sum_{(k,d,q) \in M_k^{\text{Orig}}} x_s^{(k,d,q)} \leq \sum_{\substack{j \in Out(k) \\ (k,j) \in A'}} u_s^{dc} y_{kj}^s, \quad \forall k \in \bigcup_{q \in Q} \bar{V}_{sq}^{\text{Orig}}, s \in S; \tag{25}$$

$$\sum_{(k,d,rf) \in M_k^{\text{Orig}}} x_s^{(k,d,rf)} \leq \sum_{\substack{j \in Out(k) \\ (k,j) \in A'}} u_s^{rf} y_{kj}^s, \quad \forall k \in \bar{V}_{s,rf}^{\text{Orig}}, s \in S; \tag{26}$$

$$x_s^{(o,d,q)} \leq \min\left(a^{(o,d,q)}, u_s^q\right) \sum_{\substack{j \in Out(o) \\ (o,j) \in A^{(o,d,q)}}} y_{oj}^s \quad \forall (o,d,q) \in M, s \in S; \tag{27}$$

$$x_s^{(o,d,q)} \leq \min\left(a^{(o,d,q)}, u_s^q\right) \sum_{j \in d} \sum_{\substack{i \in In(j) \\ (i,j) \in A^{(o,d,q)}}} y_{ij}^s \quad \forall (o,d,q) \in M, s \in S; \tag{28}$$

Constraints (25-26) ensure that the sum of demands loaded at any node cannot exceed the capacity of the ship, and specifically for reefer cargo. Constraints (27-28) ensure that a demand can only be carried if the ship passes through the demand's origin and one of its destinations, and caps the flow of each demand by the minimum of the demand's availability and the ship's capacity for the specific type. These constraints are the bare-basic constraints that ensure that cargo can only be carried if the ship visits the origin and a destination of the demand, and empty ships will not load more than their capacity. However, if a ship is already carrying other demands, the capacity of the ship might be breached without violating any of these constraints. This is where we add lazy constraints.

While solving a sub-problem, candidate integer solutions will be found by the MIP solver at nodes in the branch-and-bound tree. When a candidate solution is found, we check to see if the

capacity of the ship is violated at any stage, specifically at any location where additional demands are loaded. We follow the path of the ship through the network, keeping track of all demands that are currently on the ship (i.e. those that have been loaded at their origin but not yet unloaded at a destination), and if, at a particular node i, the sum of the demands currently on the ship exceeds the capacity of the ship (for any type), we then add a constraint of the form:

$$\sum_{\substack{(o,d,q)\in M \\ j\in Out(i) \\ (i,j)\in A^{(o,d,q)}}} x_s^{(o,d,q)} \leq u_s^{dc} \qquad (29) \qquad\qquad \sum_{\substack{(o,d,rf)\in M \\ j\in Out(i) \\ (i,j)\in A^{(o,d,rf)}}} x_s^{(o,d,rf)} \leq u_s^{rf} \qquad (30)$$

that is, the sum of the flow of all demands that can pass through node i must be less than the capacity of the ship, and also specifically for reefer demands. After a sub-problem has been solved, we take the lazy constraints that were generated and add them to the sub-problem as regular constraints for the next iteration. By implementing these constraints, the model is now significantly smaller and easier to solve, as can be found in Section 5.4. The results of running this implementation can be found in Section 4.4.

3.3.1. Splitting demand triples

While it's true for a specific path through the network that if two demands are carried together at any time, then they will be carried together until one is unloaded, that does not mean that the two demands must be carried together in the first place. This can occur if one of the demands has more than one possible destination. Consider the example in Figure 3.3.1, where demand A will be picked up from Origin A. There are two choices: either proceed directly to Origin B, still carrying cargo A, and pick up cargo B, or proceed to destination A1, unload cargo A and continue to Origin B. In the first case, a constraint limiting the combined capacity of demands A and B may be imposed, where it is unnecessary if the second option is chosen.

This means that the variables for demand A must be separated, which involves looking for any demand triples (o, d, q) that fit the following criteria:

- There are two destinations $d_1, d_2 \in d$ such that d_2 is reachable from d_1.
- There exists another demand m^* with origin o^* such that o^* is reachable from d_1, and d_2 is reachable from o^*.
- There exists a destination d^* of demand m^* such that d^* is reachable from d_2.
- There exists a path between the origins o and o^* that does not pass through d_1.

If such conditions are met, then the variables for $x_s^{(o,d,q)}$ are split up into $x_s^{(o,d_i,q)}$ $\forall d_i \in d$, and an additional constraint is added:

$$\sum_{d_i \in d} x_s^{(o,d_i,q)} \leq a^{(o,d,q)} \tag{31}$$

This ensures that we will not add any lazy constraints to the problem which will unnecessarily over-constrain the problem. Each demand triple has an associated revenue and amount, both of which will be inherited by the split variables.

3.3.2. Modification to the objective function

Because we are no longer explicitly calculating which arcs the cargo travels along, we cannot easily determine which destination it is being delivered to. Our implementation assumes that cargo is unloaded at the earliest possible time, that is, if the node i is visited, all demands (o, d, q) where $i \in d$ will be unloaded. In the public data set, each destination in a demand triple is the same physical port, the difference is the delivery time. This means that the unloading cost of each demand triple is the same for all destinations. This allows us to take the unload cost of any destination from the demand triple. If the destinations had different unload costs, we would have to split the demand triple as discussed in Section 3.3.1.

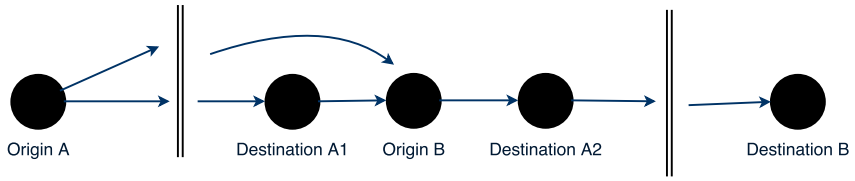

Figure 3: A scenario where the variables for a particular demand triple need to be separated.

4. Results

Here we consider the results from Tierney et. al. (2014) [20] and Tyler (2015) [23], and compare them to our results. We use the same code as Tyler, however since their results were calculated, Gurobi [10] released version 6.5 which led to significant improvements to the solve time of large MIPs. We are using a similar computer to Tyler for the reduced MIP, and report their comparison to Tierney et. al., as well as our comparison to Tyler. The machine on which the optimisation was run uses Windows 8.1 Enterprise, Python 2.7.10 and Gurobi 6.5, with an Intel i7-3770 (3.40 GHz) running 8 threads with 8 GB of RAM. All software involved is 64-bit.

Tierney et. al. tested this formulation on a public data set of 44 instances of increasing complexity. The number of ships ranged from 3 to 11, with between 30 and 379 ports and 94 to 11979 arcs. With a time limit of 1 hour and a memory limit of 10GB, the MIP implemented in CPLEX 12.4 was able to solve the first 33 instances to optimality.

4.1. Comparison of reduced model with original

Table 1 shows the results of the reduced MIP compared with the results reported in the paper by Tierney et al.. Instances are left to run for 1 hour, and if no solution is found within that time they were said to have timed out. Instances that run out of memory are denoted by 'Mem'. In the table $|M|$ represents the number of demands, $|V|$ the number of ports, $|A|$ the number of arcs and $|S|$ the number of ships.

As can be seen in Table 1 the reduced MIP formulation was able to replicate the results from Tierney et. al. for all instances. The reduced MIP appears to run much faster than that used by Tierney et. al. (2014) [20], however this is believed to be caused by an updated version of Gurobi [10] and differences in computing power, rather than differences in the model itself. Since only the smaller models are solved within the time limit, there is not much difference between the run times reported in Tyler (2015) [23] and our run times. The times achieved by our reduced MIP will serve as a baseline for comparison with the other models we present.

4.2. Revised Formulation with tighter bound

As can be seen in Table 2, the revised formulation with the tighter bound enables the MIP to solve six more instances to optimality. All solved instances are solved within 10 minutes using this formulation, which does not appear to have a significant effect on the run time of the model

Table 1: Results of reduced MIP run on public data instances from Tierney et. al. (2014) [20]

Instance ID	$\|S\|$	$\|V\|$	$\|A\|$	$\|M\|$	Tierney MIP time (s)	Tyler MIP time (s)	Reduced MIP time (s)
repos1p	3	36	150	28	0.06	0.07	0.23
repos2p	3	36	150	28	0.06	0.07	0.38
repos3p	3	38	151	24	0.04	0.06	0.25
repos4p	3	42	185	20	0.04	0.07	0.30
repos5p	3	51	270	22	0.07	0.11	0.29
repos6p	3	51	270	22	0.08	0.11	0.27
repos7p	3	54	196	46	0.08	0.09	0.26
repos10p	4	58	499	125	74.85	7.25	4.26
repos12p	4	74	603	145	106.63	14.94	11.94
repos13p	4	80	632	155	99.81	16.53	10.66
repos15p	5	71	355	173	0.47	0.29	0.40
repos16p	5	106	420	320	1.08	0.42	0.52
repos17p	6	102	1198	75	4.64	1.55	1.52
repos18p	6	135	1439	87	6.79	1.25	1.25
repos20p	6	142	1865	80	13.84	3.34	2.86
repos24p	7	75	482	154	2.23	0.47	0.58
repos25p	7	77	496	156	3.19	0.57	0.73
repos27p	7	79	571	188	1394.44	4.43	1.79
repos28p	7	90	618	189	1099.87	6.92	1.91
repos30p	8	126	1450	265	307.12	8.15	4.59
repos31p	8	130	1362	152	57.4	19.02	9.82
repos32p	8	144	1501	170	65.51	10.38	6.96
repos34p	9	304	10577	344	Time	Time	Time
repos36p	9	364	11972	1048	Mem	Time	Time
repos39p	9	379	11666	1109	Mem	Time	Time
repos41p	10	249	8051	375	Time	Time	Time
repos42p	11	279	6596	1423	Time	Time	Time
repos43p	11	320	13058	1013	Mem	Time	Time
repos44p	11	328	13705	1108	Mem	Time	Time

for smaller instances. The main benefit is that all but three of the hardest instances are solved to optimality, something which has not been achieved before. Therefore the revised formulation is a significant improvement over previous formulations for the larger instances of the LSFRP.

4.3. Column Generation

Table 2 shows the results of using column generation implemented in Python using Gurobi. Tyler's results using column generation report the approach is able to find optimal solutions to all instances in under 7.5 hrs [23]. The update to Gurobi 6.5 results in all instances being solved to optimality in under 2.7 hours. With the exception of repos43p and repos44p, all instances are solved within the timeout limit, with the slowest taking approximately 20 minutes, and the majority solving within minutes or seconds. For the larger instances (repos34p-repos44p), column generation is significantly faster than the other approaches (excluding repos41p which is slightly slower than the revised MIP). In the best case (repos34p) the column generation solves the problem almost three times faster than the revised formulation.

For the two instances that remain unsolved after the timeout limit, but can be solved within 8 hours, Tyler performs a comparison test with the revised MIP formulation [23]. This MIP is allowed to run on these instances with a new timeout limit of 8 hours, in order to determine whether it is able to solve these problems in a comparable time frame. However, after 8 hours both problems are still in the preprocessing stage of Gurobi's solver, and have not obtained a solution or an optimality gap. Therefore it can be concluded that for these large instances the only feasible technique for solving them is column generation.

4.4. Lazy constraints

All instances are solved to optimality using the lazy constraints formulation, and the optimal objective values all agreed with past results. The only comparison we are interested in is the difference in run-time between the different formulations. Table 2 is a comparison between all methods presented in this paper.

Because of the added work setting up column generation, the smaller instances for the lazy constraints method all take longer to run compared to the reduced MIP. The real strength of this formulation is that it scales far better than any other formulation presented previously. Where instances repos43p and repos44p take hours to solve with column generation, using lazy constraints

Table 2: Comparison of solution times for different implementations

Instance ID	Reduced MIP time (s)	Revised MIP time (s)	Column Generation time (s)	Lazy Constraints time (s)
repos1p	0.23	0.25	1.00	0.44
repos2p	0.38	0.28	0.53	0.50
repos3p	0.25	0.25	0.30	0.33
repos4p	0.30	0.34	0.78	0.75
repos5p	0.29	0.25	0.33	0.44
repos6p	0.27	0.30	0.34	0.34
repos7p	0.26	0.26	0.66	0.44
repos10p	4.26	1.37	3.06	1.02
repos12p	11.94	1.88	4.58	1.47
repos13p	10.66	2.02	3.75	1.65
repos14p	10.34	2.07	3.61	1.64
repos15p	0.40	0.77	1.76	1.64
repos16p	0.52	1.03	2.27	1.42
repos17p	1.52	1.85	3.06	1.80
repos18p	1.25	2.50	3.36	2.25
repos19p	1.18	2.54	3.32	2.36
repos20p	2.86	3.07	5.69	2.38
repos21p	3.14	3.13	5.64	2.70
repos22p	3.33	3.13	6.20	2.86
repos24p	0.58	1.23	2.19	1.47
repos25p	0.73	1.35	1.80	1.69
repos26p	0.63	1.34	2.10	1.38
repos27p	1.79	2.63	3.59	1.91
repos28p	1.91	2.73	3.38	2.13
repos29p	2.34	2.85	3.52	1.70
repos30p	4.59	3.09	5.83	2.78
repos31p	9.82	6.95	12.19	6.88
repos32p	6.96	3.89	6.42	3.19
repos34p	Time	251.01	91.51	19.04
repos36p	Time	493.70	228.88	38.25
repos37p	Time	447.98	263.02	33.63
repos39p	Time	539.08	374.75	35.12
repos40p	Time	541.28	370.84	29.06
repos41p	Time	88.84	113.24	17.00
repos42p	Time	Time	1271.69	57.41
repos43p	Time	Time	Time (9632.16)	223.21
repos44p	Time	Time	Time (8924.85)	222.63

Table 3: Number of ships affected by lazy constraints, and number of constraints added for each instance.

	Number of Ships		Number of Constraints	
Instance	Dry	Reef	Dry	Reef
repos10p-repos14p	1	0	4	0
repos15p-repos16p	0	3	0	7
repos36p	0	1	0	1
repos39p	0	1	0	2
repos42p	3	10	8	21
repos43p	2	10	4	22
repos44p	2	10	5	27

means the problems can be solved within four minutes. Since these larger instances are more likely to be useful in real-world situations, this formulation is by far the best for solving the LSFRP.

Another interesting result is the number of lazy constraints added during the solve process. For almost three in four instances, no lazy constraints are added at all. This means that the small formulation without flow variables for every arc is sufficient to solve the problem. The instances which most benefited from the lazy constraints formulation (repos42p-repos44p) only need at most 30 constraints for any cargo type to solve the problem to optimality. Compare this to the number of constraints that were removed from the problem, and it becomes clear why lazy constraints are so powerful when applied to this problem. A comparison of model sizes can be found in Section 5.4.

Table 3 shows, for the instances where lazy constraints are added, how many ships need them and how many constraints are added in total for the different cargo types. For all instances except the three largest, no more than seven additional constraints are required before an optimal solution was found. The reefer cargo is also more likely to have lazy constraints imposed because there are significantly fewer reefer spots on each ship compared to the dry cargo.

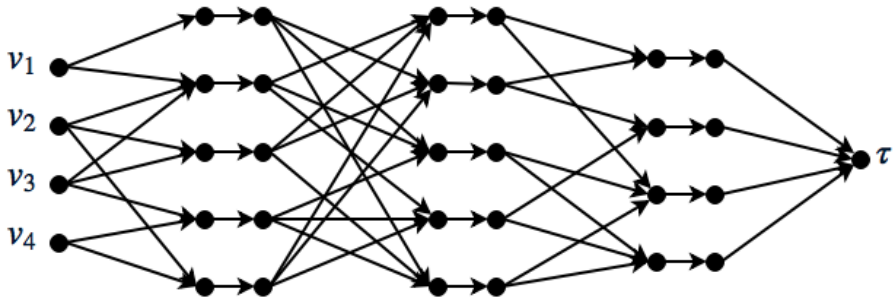

Figure 4: Example of network modification to turn column generation master problem into a pure network flow problem

5. Discussion

5.1. Proof of integer solutions to column generation

One reason column generation is so effective for this problem is because it always generates integer optimal solutions. This means we don't need to use a branch-and-bound algorithm, which makes the whole solution process much faster. After making a few observations about the nature of our problem, we will prove that solutions to the column generation are always integer.

Observation 1. *The LSFRP with a single ship type and no product flow can be modelled as a pure network flow problem.*

"The pure network flow problem can be defined by a given set of arcs and a given set of nodes with known upper bounds and cost parameters for each arc, and fixed external flow for each node" [12]. Now consider the LSFRP with a single ship type. This implies that the cost of travelling along any arc is the same for all ships. If we ignore the product flow in this problem, then we can consider the problem to have a set of source nodes v_s which each have a supply of 1 unit, a sink node τ which has a demand equal to the number of ships, and a set of intermediate nodes. These nodes can be split into In and Out nodes, with an arc connecting them which has a capacity of 1. The graphical representation of this can be seen in Figure 4. We are now solving a pure network flow problem where all constant terms are integer, which will give us the node-distinct paths of minimum cost, along which the ships will travel.

Observation 2. *Pure network flow problems have integer values at all of their extreme points*

It is known that, if the constant values of a network flow problem are all integer, then the extreme points will be integer-valued [11, 12, 8, 14].

Theorem 1. *The extreme points of the column generation master problem for a single ship type with or without cargo flows are all integer valued.*

Proof. Consider first the case without cargo flows and with one ship type. Since we can reformulate the master problem as a pure network flow problem with integer-valued constraints, the extreme points will all be integer. If we now consider the case where we have product flow, the constraints of the master problem remain unchanged. While the objective function will change, which means a different extreme point may be optimal, all extreme points are integer, so the column generation solution will still be integer. □

For Theorem 1 to apply, we require a "single ship type". This means that all ships have the same capacity and costs for all arcs. We note that this is the case for all but three instances in the public data set. In each of these instances there is one ship that is different. While it is possible that these could give non-integer solutions to the LP relaxation of the master problem, in practice they did not. To handle this possibility, we would need to modify our code to branch on the ship type that visits a node. Any continuous solution that has only one ship type visiting each node can be split into a separate solution for each ship type. Theorem 1 would then hold for each ship type, and thus the optimal solution will be integer.

5.2. Empty cargo flows in the public data set

As described in Tierney et. al. (2014) [20], the next step in this problem is to include the possibility of moving empty cargo containers for extra profit. This is a natural extension of the model proposed previously, except now instead of having a one-to-many delivery system as with the standard demands, we have a many-to-many system. This is because the capacity constraints on empty containers treat them as identical. The only distinction between the container types comes from the revenue earned, assuming one is more valuable than another. Our formulation of empty cargo flows is similar to that of the regular demands, except now we have one flow variable for each pair of supply-demand ports, where the demand port is reachable from the supply.

The parameters from Tierney et. al. (2014) [20] that are relevant to our implementation of empty cargo flows which have not previously been introduced are as follows:

$V^{q^+} \subseteq V'$	Set of visits with an equipment surplus of type q
$V^{q^-} \subseteq V'$	Set of visits with an equipment deficit of type q.
$V^{q^*} \subseteq V'$	Set of visits with an equipment surplus or deficit of type q ($V^{q^*} = V^{q^+} \cup V^{q^-}$).
$r_q^{\text{Var}} \in \mathbb{R}^+$	Revenue for each TEU of equipment of type $q \in Q$ delivered.

We denote the empty cargo flows with variables of the form $x_{sq}^{(o,d)}$, which is the amount of empty containers flowing from origin o to destination d of type q on ship s. The lazy constraints are modified appropriately, so now we keep track of which demands and how many empty containers are on board, and if the capacity is violated a constraint is added to limit these flows.

$$\sum_{\substack{(o,d,q)\in M \\ Reach(i,j,o,d,q) \\ j\in Out(i)}} x_s^{(o,d,q)} + \sum_{q\in Q} \sum_{\substack{(o,d)\in V^{q+}\times V^{q-} \\ d\in \text{ReachFrom(o)}}} x_{sq}^{(o,d)} \leq u_s^{dc} \tag{32}$$

Before presenting any results, we should note there is a problem with the public data set: it is not profitable to move empty containers. For all instances that have empty containers, the combined cost of picking up and dropping off an empty container is greater than the revenue earned from moving it. In three sets, the revenue value is 0, so even if there were no moving costs, there is no incentive to transport empty containers. Of all sets, the most profitable empty container opportunity will add -22 per container to the objective value.

As this leads to the trivial solution of $x_{sq}^{(o,d)} = 0$ for all empty cargo variables, we have decided to change the amount of revenue associated with moving empty cargo containers. All previous revenue values were either 0 or 150, but the moving costs are typically close to 150 each, so we have decided upon a new revenue value of 300. This ensures that at least moving empty containers does not result in a loss, and there are some profitable opportunities for moving empty containers, however not so profitable as to replace the movement of actual, profitable demands.

The results of this can be seen in Table 4. In some instances the increased revenue does not make a significant difference because profitable demands are already being moved, or there are no ship paths connecting (o, d) pairs. However, for the instances with more empty cargo opportunities,

Table 4: Comparison of objective value for original instances and when revenue for empty cargo fixed at 300 units.

	Without empty cargo		With empty cargo	
Instance	Objective ($\times 10^5$)	Time (s)	Objective ($\times 10^5$)	Time (s)
repos8p	-8.21	0.60	-8.21	1.04
repos9p	-8.21	0.58	-8.21	1.24
repos11p	137.61	0.88	137.61	1.94
repos14p	138.86	1.24	138.86	2.84
repos19p	5.22	2.06	15.23	4.05
repos21p	-11.85	2.17	-11.78	4.80
repos22p	-11.85	2.71	7.03	6.33
repos23p	5.22	2.09	15.06	4.14
repos26p	-53.13	1.28	-53.13	3.31
repos29p	-32.13	1.68	-32.13	4.21
repos33p	-10.92	3.32	-10.92	5.93
repos35p	138.54	26.30	157.79	114.47
repos37p	139.31	32.96	163.73	147.15
repos38p	160.02	38.47	179.27	142.00
repos40p	161.53	28.27	185.03	191.76

Table 5: Number of lazy constraints added for instances with fixed empty cargo revenue of 300

	Number of Ships		Number of Constraints	
Instance	Dry	Reef	Dry	Reef
repos19p	5	0	15	0
repos22p	6	0	21	0
repos23p	5	0	12	0
repos35p	9	1	101	1
repos37p	9	1	173	1
repos38p	9	1	94	1
repos40p	9	1	188	4

carrying empty cargo leads to higher profits. Typically these empty containers only fill previously empty spaces rather than displacing demands. The biggest difference is in the number of lazy constraints generated. Table 5 shows the instances which change when the revenue is increased. There are seven instances with empty cargo flows which do not add extra lazy constraints, however for some of the larger instances the number of constraints added approaches 200. This is still an insignificant number when compared to the size of the original formulation. All problem sets with empty cargo flows are solved using lazy constraints and column generation in less than four minutes. As such, it is no more difficult to consider the movement of empty containers in our formulation with column generation and lazy constraints.

5.3. Flexible arcs

There is one final feature of the Tierney et. al. (2014) [20] model that has not yet been implemented in our model: flexible arcs. Tierney et. al. note in their paper that the "fuel consumption of a ship is approximately a cubic function of the speed of the vessel", however in their model formulation the cost associated with sailing along a flexible arc is a linear function of sailing time, and thus has inverse relation with speed. Because the only constraints on the sailing time are upper bounds, the optimal solution for travelling along flexible arcs is always to sail as fast as possible, thus minimising the time spent sailing on the arc.

Another consideration is that flexible arcs only occur to connect flexible visits to the rest of the network. Flexible visits are ports which have no available demands, but do have a surplus or supply of empty cargo. This means that to use flexible arcs, the profit from moving empty cargo must be greater than the sum of the load and unload costs, plus the movement cost associated with the flexible arcs. For the original public data set, there are no profitable situations for moving empty cargo containers, including to flexible destinations.

5.4. Comparison of model sizes

The real benefit from lazy constraints comes from the dramatic reduction in model size, which allows the solver to manipulate the problem much more effectively. We can make a back-of-the-envelope estimate of the contribution of the demand flow variables to the size of the problem for the original formulation, and compare it to ours with column generation and lazy constraints.

In the original formulation, there are at most $|A'||M|$ variables of the form $x_{ij}^{(o,d,q)}$, where in our implementation there are less than $|M|$ variables in each sub-problem. Since there are only $|S|$

Table 6: Comparison of model sizes for instance repos44p

	Rows	Columns	Non-zeros
Reduced MIP	171102	269796	1858378
Column Generation sub-problems	136391	76612	426025
Col. Gen. and Lazy Constraints sub-problems	1627	6715	33475

sub-problems, our formulation has at most $|S||M|$ variables. For the largest instance (repos44p), $|A'| \approx 13700$, $|M| = 1108$ and $|S| = 11$, so in our formulation we have tens of thousands of variables, where the original has tens of millions.

Looking next at the number of constraints governing the x variables, the original formulation has at most $2|A'|+|M|\bullet(|V'|-1)$ constraints, where our formulation has fewer than $2|V'|+3|M|+\gamma$, where γ is the number of lazy constraints added. As seen in Table 3, $\gamma < 40$ for even the largest instance. Again looking at repos44p, where $|V'| = 327$, our formulation has fewer than 4×10^3 constraints, where the original has over 3.8×10^5. This means our model is almost three orders of magnitude smaller than the original formulation, which is why Gurobi has no problem solving the largest instances.

Table 6 is a comparison of the actual model sizes for instance repos44p between our reduced MIP, the sub-problems of the column generation implementation and of our reformulation using lazy constraints. For the latter two cases, the numbers are averaged over all 11 sub-problems. Remember that the models also include the y_{ij}^s variables. These numbers are the reported model size before Gurobi's pre-solve stage, which typically removes up to half the rows from the column generation formulation, and a few hundred rows and columns from the lazy constraints formulation. This illustrates the significant difference that using lazy constraints has made, in that the number of rows, columns and non-zeros of the column generation sub-problems have been reduced by at least one order of magnitude each.

6. Conclusion

The LSFRP is a large and important problem in the shipping industry. The only feasible way to solve large instances of the problem is with column generation, however it is the application of lazy constraints to the model which makes the most difference. By dramatically reducing the size of the model, modern solver packages can manipulate the problem much more effectively, resulting in significantly reduced run times.

This application of lazy constraints can be applied to any vehicle routing problem where transshipment cannot occur and the routes must be node-distinct. Further, it is likely that many more models can benefit from lazy constraints used in a similar way to how we have used them in the column generation sub-problems. Specifically, a large part of the model (in this case demand flow on each arc) can be optimistically approximated by a much smaller model (total amount of each demand), with additional refinements added using lazy constraints, as and when required. We consider this to be a rich area of further research.

References

[1] Agarwal, R., Ö. Ergun. 2008. Ship scheduling and network design for cargo routing in liner shipping. *Transportation Science* **42**(2) 175–196. doi:10.1287/trsc.1070.0205.

[2] Álvarez, J. 2009. Joint routing and deployment of a fleet of container vessels. *Maritime Economics & Logistics* **11**(2) 186–208. doi:10.1057/mel.2009.5.

[3] Barnhart, C., E.L. Johnson, G.L. Nemhauser, M.W.P. Savelsbergh, P.H. Vance. 1998. Branch-and-price: Column generation for solving huge integer programs. *Operations Research* **46**(3) 316–329.

[4] Brouer, B.D., J. Dirksen, D. Pisinger, C.E.M. Plum, B. Vaaben. 2013. The vessel schedule recovery problem (VSRP) a MIP model for handling disruptions in liner shipping. *European Journal of Operational Research* **224**(2) 362–374. doi:10.1016/j.ejor.2012.08.016.

[5] Brown, G.G., G.W. Graves, D. Ronen. 1987. Scheduling ocean transportation of crude oil. *Management Science* **33**(3) 335–346. doi:10.1287/mnsc.33.3.335.

[6] Christiansen, M., K. Fagerholt. 2002. Robust ship scheduling with multiple time windows. *Naval Research Logistics* **49**(6) 611–625. doi:10.1002/nav.10033.

[7] Christiansen, M., K. Fagerholt, D. Ronen. 2004. Ship routing and scheduling: Status and perspectives. *Transportation Science* **38**(1) 1–18. doi:10.1287/trsc.1030.0036.

[8] Dantzig, G. B. 1963. *Linear programming and extensions*. Princeton University Press.

[9] Fagerholt, K., M. Christiansen. 2000. A combined ship scheduling and allocation problem. *Journal of the Operational Research Society* **51** 834–842.

[10] Gurobi Optimization. 2015. *Gurobi Optimizer Reference Manual*. 6th ed. Gurobi Optimization. URL `http://www.gurobi.com/documentation/6.5/refman.pdf`.

[11] Jenson, P. 1999. Network flow programming models. URL `http://www.me.utexas.edu/~jensen/models/network/net1.html`.

[12] Jenson, P. 1999. Pure minimal cost flow problem. URL `http://www.me.utexas.edu/~jensen/methods/net.pdf/netmcf.pdf`.

[13] Kelareva, E., K. Tierney, P. Kilby. 2014. CP methods for scheduling and routing with time-dependent task costs. *EURO Journal on Computational Optimization* **2**(3) 147–194. doi:10.1007/s13675-014-0022-7.

[14] Nananukul, N. 2008. Lot-sizing and inventory routing for a production-distribution supply chain. Ph.D. thesis, University of Texas.

[15] Powell, B.J., A.N. Perkins. 1997. Fleet deployment optimization for liner shipping: An integer programming model. *Maritime Policy & Management* **24**(2) 183–192. doi:10.1080/03088839700000069.

[16] Rana, K., R.G. Vickson. 1991. Routing container ships using Lagrangean Relaxation. *Transportation Science* **25**(3) 201–214.

[17] Rathi, A.K., R.L. Church, R.S. Solanki. 1992. Allocating resources to support a multicommodity flow with time windows. *Logistics and Transportation Review* **28**(2) 167–188.

[18] Ronen, D. 1993. Ship scheduling: The last decade. *European Journal of Operational Research* **71**(3) 325–333. doi:10.1016/0377-2217(93)90343-l.

[19] Sherali, H.D., S.M. Al-Yakoob, M.M. Hassan. 1999. Fleet management models and algorithms for an oil-tanker routing and scheduling problem. *IIE Transactions* **31**(5) 395–406. doi:10.1080/07408179908969843.

[20] Tierney, K., B. Áskelsdóttir, R.M. Jensen, D. Pisinger. 2014. Solving the liner shipping fleet repositioning problem with cargo flows. *Transportation Science* 1–23doi:10.1287/trsc.2013.0515.

[21] Tierney, K., A.J. Coles, A.I. Coles, C. Kroer, A.M. Britt, R.M. Jensen. 2012. Automated planning for liner shipping fleet repositioning. *Proceedings of the 22nd International Conference on Automated Planning and Scheduling* 279–287.

[22] Tierney, K., R.M. Jensen. 2012. The liner shipping fleet repositioning problem with cargo flows. Hao Hu, Xiaoning Shi, Robert Stahlbock, Stefan Vo, eds., *Computational Logistics*, *Lecture Notes in Computer Science*, vol. 7555. Springer Berlin Heidelberg, 1–16.

[23] Tyler, A. 2015. Solving the liner shipping fleet repositioning problem with cargo flows. Unpublished honours thesis, School of Mathematics and Physics, University of Queensland.

[24] United Nations Conference on Trade and Development (UNCTAD). 2014. *Review of maritime transport*. UN Publications, New York.